

Get Latest Snowflake DEA-C02 Practice Test For Quick Preparation



P.S. Free 2026 Snowflake DEA-C02 dumps are available on Google Drive shared by BraindumpsPass:
<https://drive.google.com/open?id=1CyQeUpGdUfa1bWA6uyN6v8tl-0LA2qGF>

They are using outdated materials resulting in failure and loss of money and time. So to solve all these problems, BraindumpsPass offers actual DEA-C02 Questions to help candidates overcome all the obstacles and difficulties they face during DEA-C02 examination preparation. With vast experience in this field, BraindumpsPass always comes forward to provide its valued customers with authentic, actual, and genuine DEA-C02 Exam Dumps at an affordable cost. All the SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) questions given in the product are based on actual examination topics.

Many people dream about occupying a prominent position in the society and being successful in their career and social circle. Thus owning a valuable certificate is of paramount importance to them and passing the test DEA-C02 Certification can help them realize their goals. We treat your time as our own time, as precious as you see, so we never waste a minute or two in some useless process. Please rest assured that use, we believe that you will definitely pass the exam.

>> DEA-C02 Latest Dumps Files <<

DEA-C02 Certification Book Torrent - Latest DEA-C02 Exam Camp

As the famous saying goes, time is life. Time is so important to everyone because we have to use our limited time to do many things. Especially for candidates to take the DEA-C02 exam, time is very precious. They must grasp every minute and every second to prepare for it. From the point of view of all the candidates, our DEA-C02 training quiz give full consideration to this problem. And we can claim that if you study our DEA-C02 study materials for 20 to 30 hours, you can pass the exam for sure.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q265-Q270):

NEW QUESTION # 265

You are using the Snowflake Spark connector to update records in a Snowflake table based on data from a Spark DataFrame. The Snowflake table 'CUSTOMER' has columns 'CUSTOMER ID' (primary key), 'NAME', and 'ADDRESS'. You have a Spark DataFrame with updated 'NAME' and 'ADDRESS' values for some customers. To optimize performance and minimize data transfer, which of the following strategies can you combine with a temporary staging table to perform an efficient update?

- A. Write the Spark DataFrame to a temporary table in Snowflake. Then, execute an 'UPDATE statement in Snowflake joining the temporary table with the 'CUSTOMER table using the 'CUSTOMER_ID to update the 'NAME and 'ADDRESS' columns. Finally, drop the temporary table.
- B. Write the Spark DataFrame to a temporary table in Snowflake using MERGE. Use the WHEN MATCHED clause for Update the target table based on updates from staging table and finally drop the staging table
- C. Broadcast the Spark DataFrame to all executor nodes, then use a UDF to execute the 'UPDATE' statement for each row directly from Spark.
- D. Iterate through each row in the Spark DataFrame and execute an individual 'UPDATE statement against the

'CUSTOMER table in Snowflake. Use the 'CUSTOMER_ID in the 'WHERE clause.

- E. Use Spark's foreachPartition to batch update statements and execute on each partition. This will help with efficient data transfer and avoid single row based updates.

Answer: A,B

Explanation:

Options A and D are the most efficient. Writing to a temporary table allows Snowflake to perform the update operation using its optimized internal processes. The MERGE command is designed for efficient upserts and is preferred over individual UPDATE statements. Option B is highly inefficient due to the overhead of multiple individual queries. Option C is also not optimal, as UDFs don't necessarily improve performance for simple UPDATE operations and broadcasting data is not needed. Option E can also be an approach for batch update, its effective way for the performance but its less compared with option A and D.

NEW QUESTION # 266

- A. The Lambda function is returning a string instead of a number. Modify the Lambda function to return the discount as a number (e.g., 'discount = 0.15' instead of 'discount = '0.15")
- B. The Lambda function returns the discount within a nested JSON structure Tdata': '[[discount]]}'. The Snowflake function is not designed to handle this. The lambda function should return '{'data':
- C. The data types in the Lambda function and Snowflake function definition do not match. Specifically, the Lambda function expects strings while Snowflake is sending numbers and vice versa. Modify the Lambda function to handle numeric inputs and ensure the Snowflake function definition aligns with the expected output data type (FLOAT).
- D. The Snowflake external function is not correctly parsing the JSON response from the Lambda function. Implement a wrapper function in Snowflake to parse the JSON and extract the discount value before returning it.**
- E. The 'RETURNS NULL ON NULL INPUT clause in the external function definition is causing the function to return NULL even when valid inputs are provided. Remove this clause.

Answer: D

Explanation:

The most likely cause is (B). Snowflake expects the external function to return a single value directly convertible to the declared return type. The Lambda function is returning a JSON object that needs to be parsed. Snowflake needs a wrapper function to extract the numerical result from the json response. All other issues have been taken care of in the question and is not the cause of the problem.

NEW QUESTION # 267

You are designing a data sharing solution for a multi-tenant application where each tenant's data must be isolated. You have a 'sales' table with a 'tenant_id' column. You need to implement row-level security to ensure that each tenant can only access their own data when querying the shared table. Which of the following approaches, considering performance and security, is the MOST suitable for implementing this row-level filtering in Snowflake?

- A. Create a separate VIEW for each tenant, filtering by 'tenant_id'. Grant each tenant access only to their respective view.
- B. Implement a row access policy on the 'sales' table that filters data based on the 'tenant_id' column and the current role or user context.**
- C. Use Snowflake's data masking policies to mask all data for tenants other than the one currently querying the table.
- D. Create a scheduled task that duplicates the sales table into a new table for each tenant, filtering by the tenant_id.
- E. Implement a user-defined function (UDF) that checks the current user's tenant ID and returns a boolean value indicating whether the row should be visible. Use this UDF in a WHERE clause in every query.

Answer: B

Explanation:

Row access policies are the most efficient and scalable solution for row-level security in Snowflake. They are applied at the table level and automatically enforced for all queries, ensuring consistent data isolation. Creating separate views for each tenant is administratively burdensome and doesn't scale well. UDFs can impact performance. Data masking policies do not filter rows, they only redact data. Creating scheduled task create data duplication, its not ideal for data movements.

NEW QUESTION # 268

A data engineer is using Snowpark Python to build a data pipeline. They need to define a UDF that uses a pre-trained machine learning model stored as a file in a Snowflake stage. The UDF should receive batches of data for scoring. Which of the following is the MOST efficient way to implement this, minimizing data transfer and execution time?

- A. Use 'session.read.parquet' to load the model file directly into a Snowpark DataFrame and then use 'DataFrame.foreach' to process each row.
- B. Create a UDF that reads the model from the stage for each row that is passed to it using 'session.file.get' inside the UDF's execution logic.
- C. Use '@vectorized' decorator from Snowpark to process each batch of data passed to the UDF and load the model inside it. Specify the appropriate data types in the decorator.
- D. Load the model from the stage into a DataFrame, then use 'df.mapPartitions' to apply the model to each partition.
- E. Create a UDF with `udf(packages=['snowflake-snowpark-python', 'scikit-learn'], input_types=[ArrayType(StringType())], return_type=FloatType(), replace=True, is_peristent=True)`, and load the model within the UDF's initialization using 'session.file.get' .

Answer: C,E

Explanation:

Options C and E are the most efficient. Option C leverages the UDF framework to download the model file once when the UDF is initialized, minimizing data transfer for each call. Option E is more specific to Vectorized UDFs that takes in the batch of data at a time, so loading the model during initialization will cause less data transfer. Option A involves unnecessary DataFrame operations and partition mapping. Option B loads the entire model file into a DataFrame, which isn't necessary. Option D reads the model for each row, which is extremely inefficient.

NEW QUESTION # 269

Consider a scenario where you have a Snowflake table named 'CUSTOMER DATA' containing customer IDs (INTEGER) and encrypted credit card numbers (VARCHAR). You need to create a secure JavaScript UDF to decrypt these credit card numbers using a custom encryption key stored securely within Snowflake's internal stage, and then mask all but the last four digits of the decrypted number for data protection. Which of the following actions are necessary to ensure both functionality and security while adhering to Snowflake's best practices for UDF development and security?

- A. Encrypt the key using a weaker encryption algorithm before storing it in an internal stage to balance security and performance.
- B. Store the encryption key in a separate file on an internal stage accessible only by the UDF's service account and load the key from the file within the UDF at runtime.
- C. Use Snowflake's Secure Vault (Secret) feature to store the encryption key and retrieve it securely within the UDF.
- D. Pass the encryption key as an argument to the UDF each time it is called.
- E. Store the encryption key directly within the JavaScript UDF code as a string variable.

Answer: B,C

Explanation:

Options B and D are the correct answers. Option B - Storing the encryption key in a file on an internal stage, accessible only by the UDF's service account, is a secure way to manage the key. Option D - Snowflake's Secure Vault (Secret) feature is designed specifically for securely storing and managing sensitive information like encryption keys. This is the most recommended approach. Options A and C are insecure and should be avoided. Option E defeats the purpose of encryption.

NEW QUESTION # 270

.....

These Snowflake DEA-C02 exam questions are modeled after the DEA-C02 test. They will assist you in learning how to manage your time during the examination. BraindumpsPass enabled all users to regulate time during their SnowPro Advanced: Data Engineer (DEA-C02) DEA-C02 test. And it can be accomplished via practice, as practice makes perfect. Therefore, you must practice passing the DEA-C02 exam

DEA-C02 Certification Book Torrent: <https://www.braindumpspass.com/Snowflake/DEA-C02-practice-exam-dumps.html>

Snowflake DEA-C02 Latest Dumps Files Citing an old saying as "Opportunity always favors the ready minds", The content and

displays of the DEA-C02 pass guide Which they have tailor-designed are absolutely more superior than the other providers, If someone choose the our high efficiency exam tool, our reliable DEA-C02 dump can help users quickly analysis in the difficult point, high efficiency of review, and high quality through the exam, work for our future employment and increase the weight of the promotion, to better meet the needs of their own development, Snowflake DEA-C02 Latest Dumps Files Exam dumps are a very helpful resource during exam preparation.

However, that is not to say that these other things are inconsequential, DEA-C02 However, they are worth it, Citing an old saying as "Opportunity always favors the ready minds".

The content and displays of the DEA-C02 pass guide Which they have tailor-designed are absolutely more superior than the other providers, If someone choose the our high efficiency exam tool, our reliable DEA-C02 dump can help users quickly analysis in the difficult point, high efficiency of review, and Latest DEA-C02 Exam Camp high quality through the exam, work for our future employment and increase the weight of the promotion, to better meet the needs of their own development.

DEA-C02 Training Materials are Worthy for You to Buy It - BraindumpsPass

Exam dumps are a very helpful resource during exam preparation. Below are features of our Snowflake DEA-C02 dumps VCE file.

P.S. Free & New DEA-C02 dumps are available on Google Drive shared by BraindumpsPass: <https://drive.google.com/open?id=1CyQeUpGdUfa1bWA6uyN6v8tI-0LA2qGF>